NVM Express Technical Errata

Errata ID	**025**
Change Date	**5/24/2012**
Affected Spec Ver.	**NVM Express 1.0c**
Corrected Spec Ver.	

Submission info

Name	Company	Date
Kwok Kong	IDT	1/19/2012
Peter Onufryk	IDT	1/19/2012
Ken Okin	Virident	1/19/2012
Amber Huffman	Intel	1/19/2012

This erratum clarifies that settings for Features that are not persistent across power states is reset on a controller reset.

This erratum clarifies the LBA ranges that a host storage driver should expose; described in section 5.12.1.3. LBA ranges that are not explicitly requested to be hidden should be exposed.

Description of the specification technical flaw:

Modify the Enable field of the Controller Configuration register in section 3.1.5 as shown below:

Bit	Type	Reset	Description
00	RW	0	**Enable (EN):** When set to '1', then the controller shall process commands based on Submission Queue Tail doorbell writes. When cleared to '0', then the controller shall not process commands nor post completion queue entries to Completion Queues. When this field transitions from '1' to '0', the controller is reset (referred to as a Controller Reset). The reset deletes all I/O Submission Queues and I/O Completion Queues, resets the Admin Submission Queue and Completion Queue, and brings the hardware to an idle state. The reset does not affect PCI Express registers nor the Admin Queue registers (AQA, ASQ, or ACQ). All other controller registers defined in this section and internal controller state (e.g., Feature values defined in section 5.12.1 that are not persistent across power states) are reset to their default values. The controller shall ensure that there is no data loss for commands that have had corresponding completion queue entries posted to an I/O Completion Queue prior to the reset operation. Refer to section 7.3 for reset details. When this field is cleared to '0', the CSTS.RDY bit is cleared to '0' by the controller. When this field is set to '1', the controller sets CSTS.RDY to '1' when it is ready to process commands. The Admin Queue registers (AQA, ASQ, and ACQ) shall only be modified when EN is cleared to '0'.

Modify the first paragraph of section 5.12.1 as shown below:

Figure 73 defines the Features that may be configured with Set Features and retrieved with Get Features. Figure 74 defines Features that are specific to the NVM Command Set. Some Features utilize a memory buffer to configure or return attributes for a Feature, whereas others only utilize a Dword in the command or completion queue entry. Feature values that are not persistent across power states are reset to their default values as part of a controller reset operation. The default value for each Feature is vendor specific and is not changeable; it is set by the manufacturer.

Modify the second paragraph of section 7.3.1 as shown below:

When any of the above resets occur, the following actions are performed:

- All I/O Submission Queues are deleted.
- All I/O Completion Queues are deleted.
- All outstanding I/O commands shall be processed as aborted by host software.
- All outstanding Admin commands shall be processed as aborted by host software.
- The controller is brought to an Idle state. When this is complete, CSTS.RDY is cleared to '0'.
- All controller registers defined in section 3 and internal controller state are reset. The Admin Queue registers (AQA, ASQ, or ACQ) are not reset as part of a controller reset.

Append the following paragraph after Figure 78 as shown below:

The host storage driver should expose all LBA ranges that are not set to be hidden from the OS / EFI / BIOS in the Attributes field. All LBA ranges that follow a hidden range shall also be hidden; the host storage driver should not expose subsequent LBA ranges that follow a hidden LBA range.

Disposition log

1/19/2012	Erratum captured.
2/13/2012	Added hidden LBA range clarification.
5/10/2012	Removed read near write clarification, as this will be part of ECN 027.
5/24/2012	Clarified that default values are implementation specific.
5/24/2012	Modified default value sentence to be consistent with Technical Proposal.
7/11/2012	Erratum ratified.